

EXPRESS MAIL MAILING LABEL NUMBER EV318618171 US

**PATENT
10830.0103.NPUS00**

APPLICATION FOR UNITED STATES LETTERS PATENT

for

ON-ACCESS AND ON-DEMAND DISTRIBUTED VIRUS SCANNING

by

Luc Van Brabant

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0001] The present invention relates generally to data processing networks, and more particularly to virus scanning in a data processing network.

3. Description of Related Art

[0002] A computer virus is an intrusive program that infects computer files by inserting in those files copies of itself. When a file containing a virus is executed, the virus may replicate and infect other files, and cause damaging side effects.

9 [0003] As data networks become more open to permit a multiplicity of
10 diverse users to share access to files, the networks are subjected to an increasing threat
11 from viruses. The threat has been addressed by restricting the origin of files to trusted
12 sources, and by using virus checker software to scan the files before the files are
13 executed.

14 [0004] Virus checking software has been installed in a variety of locations
15 within a data network. Users are familiar with virus checking software that runs as a
16 background task on personal computers and workstations. This virus checking software
17 has the disadvantage that it may be disabled or in need of updating to recognize new
18 viruses.

19 [0005] Due to the relative difficulty of maintaining virus checking software
20 on workstations in a network, the virus checkers have been installed in proxy servers and
21 file servers in the network. A proxy server can function as a gatekeeper or filter for files
22 received or transmitted by a group of workstations. A proxy server having a virus
23 checker is an effective means for virus protection for services, such as electronic mail,

1 that are insensitive to transmission delay due to the time that the virus checker needs for
2 scanning the files received or transmitted. The scanning time, however, is relatively long
3 compared to the time for data access in a file server. Therefore, it is often expedient to
4 forego virus checking when accessing a file in a file server. This approach demands that
5 any file contained in the file server must have been scanned for possible viruses before
6 reading from the file. The file server, for example, contains a virus checker utility that
7 scans for viruses. When a user closes a file after any write access to the file, the file is
8 scanned for possible viruses that may have been introduced during the user's write
9 access, before any other user is permitted to open the file. If the virus checker in the file
10 server detects a virus in a file, the file remains locked by the operating system of the file
11 server until the infected file is deleted or disinfected.

12 [0006] As described in Frank S. Caccavale, U.S. Patent Application Ser.
13 09/804,320, filed March 12, 2001, entitled "Using a Virus Checker in One File Server to
14 Check for Viruses in Another File Server," Publication No. US-2002-0129277-A1,
15 incorporated herein by reference, when a network client accesses a file in a network file
16 server, the network file server invokes a conventional virus checker program in an NT
17 file server to transfer pertinent file data from the network file server to random access
18 memory in the NT file server to perform an anti-virus scan. Users may interact with the
19 virus checker program in the usual fashion, to select file types to check, and actions to
20 perform when a virus is detected. This method eliminates the need for porting the virus
21 checker program to the network file server, and avoids maintenance problems when the
22 virus checker program is updated or upgraded. Moreover, a kernel mode driver in the NT
23 file server may provide an indirect interface to the virus checker program for initiating an

1 anti-virus scan. Therefore, the driver supports a wide variety of virus checker programs
2 and ensures continued operation when the virus checker program is upgraded.

3

4 **SUMMARY OF THE INVENTION**

5 [0007] For virus checking of files in a file server, it is desirable to perform
6 “on-access” virus scanning on a priority basis, and “on-demand” virus scanning in
7 background. “On-access” virus scanning occurs when a specified trigger occurs, such a
8 when a user accesses a file marked “unchecked”. “On-demand” virus scanning is
9 typically scheduled when a new virus is discovered, when new unchecked files are
10 migrated into a file server, or prior to archiving or backing-up unchecked files.

11 [0008] Priority could be given to the “on-access” virus scan requests by
12 starting an “on-demand” virus scan only when there is no outstanding “on-access” virus
13 scan request that is not being serviced by a virus checker. In a distributed virus scanning
14 environment in which the virus checkers are separate from the file server that produces
15 the virus scan requests, however, this simple method does not keep a steady stream of
16 virus scan requests flowing to virus checkers. Conversely, it is undesirable to have too
17 many “on-demand” virus scan requests outstanding to the virus checkers, or else the
18 system will not give appropriate priority to any burst of “on-access” virus scan requests.
19 Therefore, there is a need for an improved method of allocating the “on-access” and “on-
20 demand” virus scan requests in such a distributed virus scanning environment.

21 [0009] In accordance with one aspect, the invention provides a method of
22 operating a plurality of virus checkers for on-demand anti-virus scanning concurrent with
23 on-access anti-virus scanning. The method includes combining on-demand anti-virus
24 scan requests and on-access anti-virus scan requests in a virus scan request queue, and

1 distributing the on-demand anti-virus scan requests and the on-access anti-virus scan
2 requests from the virus scan request queue to the virus checkers.

3 [00010] In accordance with another aspect, the invention provides a method of
4 operating a plurality of virus checkers. The method includes distributing on-demand
5 anti-virus scan requests and on-access anti-virus scan requests to the virus checkers so
6 that the virus checkers perform on-demand anti-virus scanning concurrent with on-access
7 anti-virus scanning. The method further includes grouping the on-demand anti-virus scan
8 requests into chunks of multiple ones of the on-demand anti-virus scan requests, and for
9 each chunk, distributing the multiple ones of the on-demand anti-virus scan requests over
10 the virus checkers.

11 [00011] In accordance with yet another aspect, the invention provides a method
12 of operating a plurality of virus checkers for on-demand anti-virus scanning concurrent
13 with on-access anti-virus scanning. The method includes combining on-demand anti-
14 virus scan requests and on-access anti-virus scan requests in a virus scan request queue,
15 and a pool of threads distributing the on-demand anti-virus scan requests and the on-
16 access anti-virus scan requests from the virus scan request queue to the virus checkers.
17 Each anti-virus scan request on the virus scan request queue is serviced by a respective
18 one of the threads in the pool of threads. The method further includes grouping the on-
19 demand anti-virus scan requests into chunks of multiple ones of the on-demand anti-virus
20 scan requests, and for each chunk, checking whether the number of anti-virus scan
21 requests on the virus checking queue is less than a threshold, and upon finding that the
22 number of anti-virus scan requests on the virus checking queue is less than the threshold,
23 placing the chunk on the virus scan request queue.

1

2 **[00012]** In accordance with still another aspect, the invention provides a virus
3 checking system including a plurality of virus checkers for on-demand anti-virus
4 scanning concurrent with on-access anti-virus scanning, a virus scan request queue, and
5 at least one processor coupled to the virus checkers and the virus scan request queue for
6 sending virus scan requests from the virus scan request queue to the virus checkers. The
7 at least one processor is programmed for placing on-demand anti-virus scan requests and
8 on-access anti-virus scan requests onto the virus scan request queue, and for distributing
9 the on-demand anti-virus scan requests and the on-access virus scan requests from the
10 virus scan request queue to the virus checkers.

11 **[00013]** In accordance with a final aspect, the invention provides a virus
12 checking system including a plurality of virus checkers for on-demand anti-virus
13 scanning concurrent with on-access anti-virus scanning, and a file server coupled to the
14 virus checkers for sending virus checking requests from the file server to the virus
15 checkers. The file server includes a virus scan request queue. The file server is
16 programmed for placing on-demand anti-virus scan requests and on-access anti-virus
17 scan requests onto the virus scan request queue, and for executing multiple threads for
18 distributing the on-demand anti-virus scan requests and the on-access anti-virus scan
19 requests from the virus scan request queue to the virus checkers. Each anti-virus scan
20 request on the virus scan request queue is serviced by a respective one of the threads in
21 the pool of threads. The file server is further programmed for grouping the on-demand
22 anti-virus scan requests into chunks of multiple ones of the on-demand anti-virus scan
23 requests, and for consecutively placing the chunks onto the virus scan request queue.

BRIEF DESCRIPTION OF THE DRAWINGS

[00014] Other objects and advantages of the invention will become apparent upon reading the detailed description with reference to the drawings, in which:

[00015] FIG. 1 is a block diagram of a data processing system incorporating the present invention;

[00016] FIG. 2 is a flowchart of a method of using a virus checker in the system of FIG. 1;

[00017] FIG. 3 is a block diagram showing programming and data structures in a data mover in the data processing system of FIG. 1;

[00018] FIGS. 4 to 6 comprise a flowchart of an antivirus (AV) thread in the data mover of FIG. 3;

[00019] FIG. 7 is a flowchart of an on-demand request chunking thread in the data mover of FIG. 3;

[00020] FIG. 8 is a flowchart of a subroutine used in FIG. 7 for dumping a chunk of on-demand virus checking requests onto the virus scan request queue; and

[00021] FIG. 9 is a flowchart of a procedure for virus checking system configuration and adjustment.

18 [00022] While the invention is susceptible to various modifications and
19 alternative forms, a specific embodiment thereof has been shown by way of example in
20 the drawings and will be described in detail. It should be understood, however, that it is
21 not intended to limit the form of the invention to the particular form shown, but on the
22 contrary, the intention is to cover all modifications, equivalents, and alternatives falling
23 within the scope of the invention as defined by the appended claims.

1

2 **DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS**

3 **[00023]** With reference to FIG. 1, there is shown a distributed data processing
4 system incorporating the present invention. The data processing system includes a data
5 network 21 interconnecting a number of clients and servers. The data network 21 may
6 include any one or more of network connection technologies, such as Ethernet or Fibre
7 Channel, and communication protocols, such as TCP/IP or UDP. The clients include
8 work stations 22 and 23. The work stations, for example, are personal computers. The
9 servers include conventional Windows NT/2000 file servers 24, 25, 26, and a very large
10 capacity network file server 27. The network file server 27 functions as a primary server
11 storing files in nonvolatile memory. The NT file servers 24, 25, 26 serve as secondary
12 servers performing virus checking upon file data obtained from the network file server
13 27. The network file server 27 is further described in Vahalia et al., U.S. Patent
14 5,893,140 issued April 6, 1999, incorporated herein by reference. Such a very large
15 capacity network file server 27 is manufactured and sold by EMC Corporation, 176 South
16 Street, Hopkinton, Mass. 01748.

17 **[00024]** The network file server 27 includes a cached disk array 28 and a
18 number of data movers 29, 30 and 31. The network file server 27 is managed as a
19 dedicated network appliance, integrated with popular network operating systems in a
20 way, which, other than its superior performance, is transparent to the end user. The
21 clustering of the data movers 29, 30, 31 as a front end to the cached disk array 28
22 provides parallelism and scalability. Each of the data movers 29, 30, 31 is a high-end
23 commodity computer, providing the highest performance appropriate for a data mover at

1 the lowest cost. The network file server 27 also has a control station 35 enabling a
2 system administrator 36 to configure and control the file server.

3 [00025] Each of the NT file servers 24, 25, 26 is programmed with a respective
4 conventional virus checker 32, 33, 34. The virus checkers are enterprise class anti-virus
5 engines, such as the NAI/McAfee's NetShield 4.5 for NT Server, Symantec Norton
6 AntiVirus 7.5 Corporate Edition for Windows NT, Trend Micro's ServerProtect 5.5 for
7 Windows NT Server. In each of the NT file servers 24, 25, 26, the virus checker 32, 33,
8 34 is invoked to scan a file in the file server in response to certain file access operations.
9 For example, when the file is opened for a user, the file is scanned prior to user access,
10 and when the file is closed, the file is scanned before permitting any other user to access
11 the file.

12 [00026] The network file server 27, however, is not programmed with a
13 conventional virus checker, because a conventional virus checker needs to run in the
14 environment of a conventional operating system. Network administrators, who are the
15 purchasers of the file servers, would like the network file server 27 to have a virus
16 checking capability similar to the virus checking provided in the conventional NT file
17 servers 24, 25, 26. Although a conventional virus checker could be modified to run in the
18 environment of the data mover operating system, or the data mover operating system
19 could be modified to support a conventional virus checker, it is advantageous for the
20 network file server 27 to use the virus checkers 27, 28, 29 in the NT file servers to check
21 files in the network file server 27 in response to user access of the files in the network file
22 server. This avoids the difficulties of porting a conventional virus checker to the network
23 file server, and maintaining a conventional virus checker in the data mover environment

1 of the network file server. Moreover, in many cases, the high-capacity network file
2 server 27 is added to an existing data processing system that already includes one or more
3 NT file servers including conventional virus checkers. In such a system, all of the files in
4 the NT file servers 24, 25, 26 can be migrated to the high-capacity network file server 27
5 in order to facilitate storage management. The NT file servers 24, 25, 26 in effect
6 become obsolete for data storage, yet they can still serve a useful function by providing
7 virus checking services to the network file server 27.

8 [00027] In general, when a client 22, 23 stores or modifies a file in the network
9 file server 27, the network file server determines when the file needs to be scanned.
10 When anti-virus scanning of a file has begun, other clients are blocked on any access to
11 that file, until the scan completes on the file. The network file server 27 selects a
12 particular one of the NT file servers 24, 25, 26 to perform the scan, in order to balance
13 loading upon the NT file servers for anti-virus scanning processes. The virus checker in
14 the selected NT file server performs a read-only access of the file to transfer file data
15 from the network file server to random access memory in the selected NT file server in
16 order to perform the anti-virus scan in the NT file server. Further details regarding the
17 construction and operation of the virus checkers 32, 33, 34 and the interface between the
18 virus checkers and the network file server 27 are found in Caccavale United States Patent
19 Application Publication No. US 2002/0129277 A1 published Sep. 12, 2002, incorporated
20 herein by reference.

21 [00028] With reference to FIG. 2, there is shown a flowchart of a method of
22 using a virus checker in the system of FIG. 1. In a first step 50 of FIG. 2, a client (22, 23
23 in FIG. 1) sends new data to the primary network file server (27 in FIG. 1). Next, in step

1 51, the primary network file server receives the new data, and selects one of the virus
2 checkers for load balancing. For example, a virus checker is selected using a “round
3 robin” method that places substantially the same workload upon each of the virus
4 checkers. For the case of the three NT file servers 24, 25, and 26 in FIG. 1, for example,
5 a distribution list [NT1, NT2, NT3] is stored in each data mover, and each time that a
6 virus scan request from a data mover needs to be serviced, the data mover increments a
7 pointer to the distribution list to distribute the scanning requests to the NT file servers in a
8 sequence [NT1, NT2, NT3, NT1, NT2, NT#, NT1, NT2, NT3, ...], where NT1 is the
9 name of the first NT file server 24, NT2 is the name of the second NT file sever 25, and
10 NT3 is the name of the third NT file server 26. In general, for the case of a total number
11 N of active virus checkers, and a total workload of W, the load upon each active virus
12 checker is W/N .

13 **[00029]** In step 52, the primary network file server sends an anti-virus scan
14 request to the NT file server (24, 25, or 26 in FIG. 1) having the selected virus checker
15 (32, 33, or 34). The scan request identifies the new data to be scanned. In step 53, the
16 selected virus checker receives the scan request, and accesses the new data in the primary
17 network file server. In step 54, the selected virus checker determines if there is a risk of a virus
18 being present in the new data, and recommends an action if there is a risk of a virus
19 being present. The virus checking for the new data is finished after step 54.

20 **[00030]** The scanning task shown in FIG. 2 is generally referred to as an “on-
21 access” virus scan. An “on-access” virus scan is processed in real-time when scanning is
22 triggered by user-initiated file access. Another kind of virus scan is known as an “on-
23 demand” virus scan. An “on-demand” virus scan is scheduled at a lower priority than

1 “on-access,” and it typically involves scanning all files of virus-checkable file type in a
2 one or more specified file systems. For example, “on-demand” virus checking is
3 scheduled when a new virus is discovered, when new unchecked files are migrated into a
4 file server, or prior to archiving or backing-up unchecked files. Although lower in
5 priority from a scheduling point of view, “on-demand” virus checking has often been
6 more burdensome on the data processing system than “on-access” virus checking. A full
7 file system scan may generate a much more intense scanning load when a multitude of
8 files in the file system must be scanned. Even though this scanning workload is
9 distributed over multiple virus checkers, the volume of scans will generate a significant
10 resource load on the operating system of the data mover. Moreover, it is desirable to
11 fully utilize the capabilities of the virus checkers in order to complete the full file system
12 scan as soon as possible

13 **[00031]** In order to mitigate any general performance degradation on the data
14 mover during user file access, it is desirable to mix “on-demand” virus scan requests with
15 “on-access” virus scan requests in a shared virus scan request queue. For example,
16 outstanding “on-demand” virus scan requests are added to the shared queue when the
17 number of requests in the shared queue falls below a threshold. The threshold is selected
18 to provide a relatively continuous flow of requests to the virus checkers without
19 significantly degrading the response time of the virus checkers for responding to the “on-
20 access” requests. Moreover, it is desirable to add outstanding “on-demand” virus scan
21 requests to the shared queue in manageable “chunks”, and to wait until the virus scan
22 requests in each chunk have been serviced before sending another chunk of “on-demand”
23 virus scan requests.

1
2 **[00032]** As shown in FIG. 3, the data mover 29 stores a database 61 of virus
3 checkers, and a distribution list 62 of the virus checkers that are in a condition to service
4 virus scan requests. Virus scan requests to be serviced are listed in a virus scan request
5 queue 63. The virus scan request queue 63 is serviced by a pool of anti-virus threads 64.
6 Anti-virus threads that are presently servicing virus scan requests are listed in a queue 65
7 of active anti-virus threads 65.

8 **[00033]** On-access virus scan requests and on-demand virus scan requests are
9 placed in the virus scan request queue 63. When an on-access virus scan request occurs
10 due to user access to a file that needs to be scanned, the on-access virus scan request is
11 immediately placed on the virus scan request queue 63. When an on-demand virus scan
12 request is received from the system administrator (36 in FIG. 1), the request is placed on
13 an on-demand request queue 66. An on-demand request chunking routine 67 services the
14 on-demand request queue 66 by grouping on-demand anti-virus file scan requests into
15 chunks, and when the number of anti-virus file scan requests in the virus scan request
16 queue is less than a threshold, by inserting each chunk of on-demand anti-virus file scan
17 requests into the virus scan request queue 63.

18 **[00034]** For example, there are ten anti-virus threads in the anti-virus thread
19 pool 64, there are five virus checkers, and the chunk size is fifty anti-virus file scan
20 requests. The threshold is set to twenty-five. In a more general case of “N” virus
21 checkers and “M” anti-virus threads, for example, the chunk size is the product (M*N),
22 and the threshold is set to one-half of the chunk size.

23 **[00035]** FIGS. 4 to 6 show a flowchart of an antivirus (AV) thread in the data
24 mover of FIG. 3. In a first step 71, the AV thread obtains a lock on the head of the virus

1 scan request queue. In step 72, if the virus scan request queue is empty, then execution
2 branches to step 73, to release the lock. In step 74, execution of the AV thread is
3 suspended for a time, and then execution resumes to loop back to step 71.

4 [00036] In step 72, if the virus scan request queue is not empty, then execution
5 continues to step 75. In step 75, the AV thread is placed on the tail of the queue of active
6 AV threads. In step 76, a file path name is removed from the head of the virus scan
7 request queue. This is the name of the file to be serviced by the AV thread. In step 77,
8 the AV thread releases the lock on the head of the virus scan request queue.

9 [00037] In step 78, the AV thread gets a lock on the distribution list. In step
10 79, the AV thread increments the distribution list pointer. In step 80, the AV thread
11 accesses the distribution list with the pointer to get the virus checker that will scan the file
12 being serviced by the AV thread. In step 81, the AV thread releases the lock on the
13 distribution list. In step 82, the AV thread accesses the database of virus checkers to get
14 the URL of the virus checker assigned to scan the file being serviced by the AV thread.
15 Execution continues from step 82 to step 83 in FIG. 5.

16 [00038] In step 83 of FIG. 5, the AV thread sends the virus scan request to the
17 URL of the assigned virus checker. In step 84, the AV thread records the present time
18 when the virus scan request was sent to the assigned virus checker. In step 85, the AV
19 thread calls a time service to set up a callback to the virus checker after a timeout
20 interval. In step 86, the AV thread is suspended while waiting for a callback from the
21 assigned virus checker or from the timer service.

22 [00039] In step 87, if the AV thread receives a callback from the assigned virus
23 checker, then execution continues to step 88. In step 88, the AV thread cancels the timer

1 service. In step 89, the AV thread looks at a return code from the virus checker. If the
2 return code indicates an error condition, then execution continues from step 89 to step 90.
3 In step 90, the AV thread forwards the error code and the file pathname to an error
4 handler. For example, the error code may indicate that a virus has been found in the
5 file, in which case the error handler will log the error and deny user access to the file.
6 Execution continues from step 90 to step 91. Execution also continues to step 91 from
7 step 89 if the virus checker did not return an error.

8 [00040] In step 91, the AV thread sends a callback to the process that requested
9 the virus scan of the file, and the AV thread is removed from the queue of active AV
10 threads. Execution continues from step 91 to step 92 in FIG. 4. In step 92, the AV thread
11 tests a flag to see if the control station has disabled the AV thread. If so, then the AV
12 thread terminates. Otherwise, execution loops from step 92 back to step 71.

13 [00041] In step 87 in FIG. 5, if the callback was not from the assigned virus
14 checker, then execution continues to step 93. In step 93, the callback should be from the
15 timer service. If not, then execution loops back to step 86. Otherwise, execution
16 continues from step 93 to step 94 in FIG. 6. In step 94, the AV thread reports the time-
17 out error to the control station. The AV thread assumes that the virus checker is
18 defective, and therefore, in steps 95 to 98, the virus checker re-assigns the virus checking
19 task to another virus checker. In step 95, the AV thread gets a lock on the distribution
20 list. In step 96, the AV thread increments the distribution list pointer. In step 97, the AV
21 thread accesses the distribution list with the pointer to get another virus checker to service
22 the virus scan request. In step 98, if the pointer does not point to a virus checker that is
23 different from the virus checker that had been assigned to the request, then execution

1 loops back to step 96 to assign the next virus checker in the distribution list. Once a
2 different virus checker has been assigned, execution loops back from step 98 to step 81 in
3 FIG. 4.

4 [00042] FIG. 7 shows a flowchart of the routine for on-demand request
5 chunking. This routine services the on-demand request queue (66 in FIG. 3). For each
6 file system listed in the on-demand request queue, the routine finds the files that should
7 be scanned for viruses, and groups these files into chunks that are consecutively placed
8 on the virus scan request queue (66 in FIG. 3).

9 [00043] In a first step 100 of FIG. 7, a list of files, called the chunk, is cleared.
10 Next, in step 101, the chunking routine accesses the on-demand request queue. If the
11 queue is not empty, then execution continues to step 103. In step 103, the chunking
12 routine accesses the root directory of the file system at the head of the on-demand request
13 queue. In step 104, the chunking routine begins a search for all of the files in the file
14 system. In step 105, if the end of the file system has not been reached, then execution
15 continues to step 106. In step 106, the chunking routine gets the next file in the file
16 system. In step 107, if the file type is such that the file should not be scanned (e.g., it is
17 not an executable file), then execution loops back to step 105. Otherwise, execution
18 continues from step 107 to step 108. In step 108, the file is added to the chunk. In step
19 109, if the size of the chunk has not reached the maximum number of files for the chunk,
20 then execution loops back to step 105. Otherwise, execution branches to step 110. In
21 step 110, a subroutine shown in FIG. 8 dumps the chunk onto the tail of the virus scan
22 request queue when it is appropriate to do so. After step 110, execution loops back to
23 step 105.

1 **[00044]** In step 105, once the end of the file system has been reached, then
2 execution branches to step 113. In step 113, the request chunking routine removes the
3 entry at the head of the one-demand request queue, and execution loops back to step 101.

4 **[00045]** In step 102, once the on-demand request queue is empty, execution
5 branches to step 114. In step 114, if the chunk is empty, then execution continues to step
6 115 to suspend and resume the request chunking routine. From step 115, execution loops
7 back to step 101.

8 **[00046]** In step 114, if the chunk is not empty, then execution continues to step
9 116. In step 116, the subroutine of FIG. 8 dumps the chunk onto the tail of the virus-
10 scan request queue when it is appropriate to do so. From step 116, execution continues to
11 step 115.

12 **[00047]** FIG. 8 shows the subroutine referenced in steps 110 and 116 of FIG. 7
13 for dumping the chunk onto the tail of the virus scan request queue when it is appropriate
14 to do so. In a first step 121 of FIG. 8, execution continues to step 122 if a prior chunk has
15 been put on the virus scan request queue. In step 121, if all of the files in the prior chunk
16 have not been scanned, then execution branches to step 123. In step 123, execution is
17 suspended and resumes, and execution loops back to step 122. Once all of the files in the
18 prior chunk has been scanned, execution continues from step 122 to step 124. Execution
19 also branches from step 121 to step 124 if a prior chunk has not been put on the virus
20 scan request queue.

21 **[00048]** In step 124, if the number of entries in the virus scan request (VSR)
22 queue is not less than the threshold (TH1), then execution continues to step 125 to
23 suspend and resume execution. After step 125, execution loops back to step 124. Once

1 the number of entries in the virus scan request queue is less than the threshold (TH1),
2 execution continues to step 126. In step 126, the list of file names from the chunk are
3 dumped onto the tail of the virus scan request queue, and execution returns.

4 [00049] FIG. 9 shows a flowchart of a procedure for configuration and
5 adjustment of the virus checking system. In a first step 131, the number (M) of virus
6 checking threads is set at a multiple (n) of the number (N) of virus checkers in order to
7 have up to the multiple (n) of outstanding virus checking requests distributed to each of
8 the virus checkers. In step 132, the number (M) of virus checking threads can be
9 readjusted when the virus checking system is online. For example, a flag (tested in step
10 92 of FIG. 4) can be set for each active virus checker to terminate the active virus checker
11 once it has finished with servicing of a current request. In step 133, the queue threshold
12 (TH1) and the chunk size can be adjusted to set the relative priority of the on-demand
13 requests relative to the on-access requests. A near-empty queue threshold (TH1) and a
14 relatively small chunk size will result in the on-demand requests being placed entirely in
15 background during high loading conditions. However, a larger chunk size and a larger
16 queue threshold will tend to keep the virus checkers busy under diverse loading
17 conditions.

18 [00050] In view of the above, on-demand and on-access virus scan requests are
19 distributed over a plurality of virus checkers for on-demand anti-virus scanning
20 concurrent with on-access anti-virus scanning. This method reduces the scan time for on-
21 demand scanning of a file system and eliminates a potential single point of failure since
22 more than one virus checker is used for scanning of the file system. The method can also
23 keep multiple virus checkers busy scanning the files in the file system without

1 substantially reducing the availability of the virus checkers for on-access virus checking.
2 For example, the on-demand anti-virus scan requests are grouped into chunks of multiple
3 requests, and the on-demand and on-access requests are combined in a queue. Scanning
4 for requests from a prior chunk are completed before distributing the requests for a next
5 chunk. To give priority to the on-access requests, the on-demand requests are not placed
6 on the queue unless the number of requests on the queue is less than a threshold. An on-
7 access request for anti-virus scanning of a file is placed on the queue in response to a
8 request for user access to the file. On-demand requests are produced in response to a
9 request from a system administrator for anti-virus scanning of the files in a specified file
10 system.